Fast Approximate Matching of Cell-Phone Videos for Robust Background Subtraction

Raffay Hamid, Atish Das Sarma, Dennis DeCoste, Neel Sundaresan
raffay@gmail.com, atish.dassarma@gmail.com, ddecoste@ebay.com, nsundaresan@ebay.com
eBay Research Laboratory, San Jose, CA. USA.

arXiv:1404.5351v1 [cs.CV] 22 Apr 2014

ABSTRACT

We identify a novel instance of the background subtraction problem that focuses on extracting near-field foreground objects captured using handheld cameras. Given two user-generated videos of a scene, one with and the other without the foreground object(s), our goal is to efficiently generate an output video with only the foreground object(s) present in it. We cast this challenge as a spatio-temporal frame matching problem, and propose an efficient solution for it that exploits the temporal smoothness of the video sequences. We present theoretical analyses for the error bounds of our approach, and validate our findings using a detailed set of simulation experiments. Finally, we present the results of our approach tested on multiple real videos captured using handheld cameras, and compare them to several alternate foreground extraction approaches.

1. INTRODUCTION

Over the years, there has been a tremendous increase in the number of videos recorded from cell-phones and other handheld mobile devices. A majority of these videos capture a set of near-field objects in visually cluttered environments. Automatic extraction of these foreground objects in such videos is an important problem, with potential impacts on a variety of different application areas.

For videos captured using handheld cameras, it is relatively easy to record a few frames of the scene without the foreground objects present, to coarsely characterize what the background looks like. We argue that this extra, albeit approximate, information about the background scene can be very useful to extract foreground objects in such videos.

The motivation of our novel problem setting is grounded in several real-world applications. In online commerce for instance, thousands of sellers use mobile phones to capture videos of their products they want to sell online. Most of these sellers are amateur videographers who lack the skills or the equipment to capture good quality videos of their products. An important factor adversely affecting their video quality is the background scene clutter. As improving the quality of product videos increases the likelihood of selling the products, these users are naturally incentivized to provide a short additional video of the background scene in return for having product videos with improved quality.

Similarly on filming rigs, it is a standard practice to capture background and foreground scene videos along carefully planned camera paths using expensive computerized dollies to generate various special effects. By enabling robust and efficient background-foreground video matching using only hand-held cameras, our work attempts to make this process more efficient and less expensive. Furthermore, our work attempts to open up avenues for amateur videographers who cannot afford costly gear, to generate these special effects using only their handheld mobile cameras.

The precise objective of our work is, given two handheld videos, one with and the other without foreground object(s), to efficiently generate a video with only the foreground object(s) shown. Note that if camera paths for both the foreground and the background videos in our problem were identical, we could simply perform a pixel-wise background subtraction. On the other hand, if the two camera paths were entirely different with no scene overlap, the background video would be redundant and our problem would transform to an object tracking challenge. In practice however, one is somewhere in the middle of this spectrum with some amount of partial path overlap between the two videos. The challenge therefore is, for each foreground frame, to find a set of background frames with sufficient scene overlap, such that with the appropriate frame-alignment, they could be used to subtract out the background pixels from the foreground frame.

A naive approach to this spatio-temporal frame matching problem requires operations quadratic in the number of the foreground and background frames. For only one minute long background and foreground videos captured at 30 frames per second, a naive approach requires more than 3 million pair-wise frame comparisons. Even with an optimized implementation exploiting heterogeneous compute architecture, each comparison can take up to 100 ms, which would result in the entire process taking several days to complete. It is therefore crucial that for each foreground frame, we find the relevant background frames more efficiently.

The **main contributions** of our work are:

• We propose an approximation algorithm for efficient spatio-temporal matching of videos captured using hand-held cameras. Our algorithm permits a significantly improved computational complexity compared to an exhaustive approach.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

• We present a systematic way to quantify bounds on the error incurred by our approximate matching algorithm, and corroborate our theoretical findings by a set of controlled simulation experiments.

• We present a computational framework using our proposed frame matching algorithm to efficiently subtract the scene background given a background and a foreground video.

• We validate the robustness of our framework on numerous real videos, and show the improved performance of our technique compared to multiple alternate approaches.

In the following, we start by going over the relevant previous works identifying how our work is different from them. We then concretely formalizing our spatio-temporal frame matching problem. We use the insights from our formulation to make appropriate choices for our computational framework explained in § 5.

2. RELATED WORK

Foreground extraction in videos has previously been explored in great detail [25] [28], however most of this work has either focused on static camera(s), or for cameras with limited motion [14] [4]. To overcome this challenge, there have been several efforts to look at the problem purely from a tracking based perspective [15] [7]. While these approaches can successfully localize object position, they are not geared for accurately extracting object silhouettes.

A solution for this challenge is to use tracking and segmentation in closed-loop form [26] [20]. However, such approaches either require human intervention to provide object silhouette, or assume specific lighting conditions.

Another method to estimate object silhouette to initialize tracking is to segment the scene into multiple motion layers [23] [27]. However, such methods are limited to non-static scenes with distinct object-level motion boundaries.

There have been some recent attempts to combine appearance and stereo cues for foreground extraction. Most of these methods however perform full sequence calibration [17], or assume that the background scene can be approximated using piecewise planar geometry [11]. There have also been attempts to extract foreground object(s) based on frame saliency [16] [12] and co-segmentation [19].

Recently there has been work on utilizing the general background appearance sampled sparsely over space and time from unregistered images [2] [1]. We show in § 6.3.1 however that this model does not generalize for handheld video cameras as their unconstrained motion requires quite a dense sampling of background scene appearance, resulting in prohibitively high capture and computational cost. Similar approaches for scene alignment have been explored for videos [21] [5] [6]. However, unlike our work, they assume same or near identical camera paths.

3. PROBLEM SETUP

Consider Figure 1, where a 3-D scene is shown to be captured by a handheld camera tracing two different paths. For one of these captures, the scene contains a (set of) foreground object(s) resulting in a foreground video. For the second capture, a background video is recorded without the foreground object(s). Let n and m denote the number of foreground and background video frames respectively. Let $\bar{f} = \{f_1, f_2, \ldots, f_n\}$ and $\bar{b} = \{b_1, b_2, \ldots, b_m\}$ denote the foreground and background frame sequences respectively. We use $d(x, y)$ to denote the distance between frames x and y, the exact notion of which is formally defined in § 5.1.

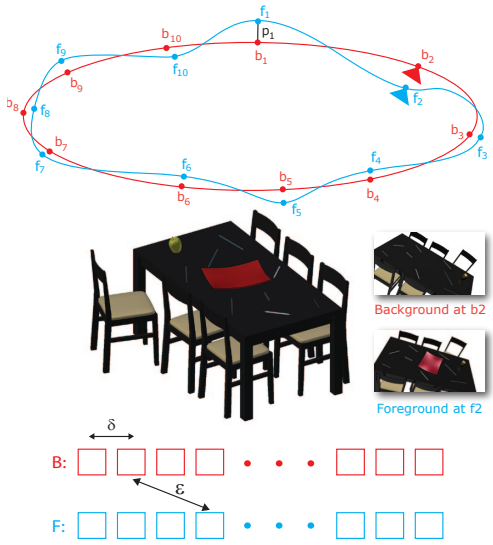

Figure 1: An example 3-D scene is captured by a handheld camera tracing two different paths, f and b, one with and the other without a foreground object (red tray). This results in a foreground and a background video, $\bar{f}$ and $\bar{b}$ respectively. Example frames from these videos for this scene are also shown at f_2 and b_2.

3.1 Problem Assumptions

We make the following assumptions to setup our problem.

3.1.1 *Path Smoothness*

We assume that consecutive frames in a video are similar to each other, *i.e.*, the path traced by the camera during a capture is generally *smooth*. Given the high capture rate of modern mobile devices, this assumption is well-justified. Formally, $\forall\ i\ (1 \leq i < n)$, we have $d(f_i, f_{i+1}) \leq \delta$, and $\forall\ j\ (1 \leq j < m)$, we have $d(b_j, b_{j+1}) \leq \delta$, for $\delta > 0$.

3.1.2 *Capture Completeness*

Given our problem setup, it is likely that several foreground frames would have strong background matches, *i.e.*, for such foreground frames the background capture would be *complete*. For the sake of clarity, we begin with the assumption that capture completeness applies $\forall f \in \bar{f}$. We address the scenario when this is not the case in § 5.6. Formally, for each i between $1 \leq i \leq n$, $\exists\ j$ between $1 \leq j \leq m$ such that $d(f_i, b_j) \leq \epsilon$. Note that here ϵ is significantly larger than δ (as the object of interest is present only in the foreground video), yet typically small when the background video is fairly exhaustive.

3.2 Problem Formulation

Given $\bar{f}$ and $\bar{b}$, $\forall\ 1 \leq i \leq n$, we want to find a corresponding $1 \leq j \leq m$ to minimize $d(f_i, b_j)$. More formally, we want to identify function $\pi : [1, n] \rightarrow [1, m]$ such that $\texttt{avg}_{i=1}^{i=n} d(f_i, b_{\pi(i)})$ is minimized. Note that here we use the

Algorithm 1 NEAR-LINEAR(k)

Input: $\bar{f}$, $\bar{b}$, and parameter k.
Output: π, matched backgrounds $\forall f \in \bar{f}$.

1: Let $S_i^k = \{i : 1 \leq i \leq n, i = 0 \mod k\} \cup \{1\} \cup \{n\}$.
2: Let $S_j^k = \{j : 1 \leq j \leq m, j = 0 \mod k\} \cup \{1\} \cup \{m\}$.
3: $\forall i \in S_i^k$, set $\pi(i) = \arg\min_{j \in S_j^k} d(f_i, b_j)$.
4: $\forall i \notin S_i^k$, let $near(i) = \arg\min_{t \in S_i^k} |t - i|$.
5: $\forall i \notin S_i^k$, set $\pi(i) = \pi(near(i))$.

Algorithm 2 LOCAL-SEARCH($\bar{f}, \bar{b}$)

Input: f, $\bar{b}$, k, and γ.
Output: Matched background $\forall f \in \bar{f}$.

1: Let $S_j^k = \{j : 1 \leq j \leq m, j = 0 \mod k\} \cup \{1\} \cup \{m\}$.
2: For $j \in S_j^k$, compute if b_j is a *strong match* for f.
3: $\forall\, b_j$ that are strong matches, compare f with each $b \in b_x$ where $x : [j - \gamma, j + \gamma]$ to find other strong matches.
4: Return all strong matches found.

average cost; however one could define the objective using max (or equivalent functions) and similar results hold.

4. SPATIO-TEMPORAL MATCHING

In the following, we first present a naive algorithm for frame matching between the foreground and background videos, followed by our proposed approximate algorithm that guarantees a significant speedup while only incurring minimal accuracy cost. For all the proofs of theorems and lemmas, we refer the reader to Appendix A.

4.1 Naive Algorithm

Recall that matching is determined by a function π from $[1, n]$ to $[1, m]$, and that the cost of *any* matching, π, is given by $C(\pi) = \text{avg}_{i=1}^{i=n} d(f_i, b_{\pi(i)})$. The optimal matching is one that minimizes this cost. Let us denote the optimal matching by π^* and the associated cost by $C(\pi^*)$. That is, $\pi^* = \arg\min_\pi \text{avg}_{i=1}^{i=n} d(f_i, b_{\pi(i)})$. The path smoothness assumption results in the following lemma.

Lemma 4.1. $C(\pi^*) \leq \epsilon$.

A naive brute-force matching algorithm finds the best background match for every foreground frame, *i.e.*, for each f_i, compute $d(f_i, b_j)\ \forall\, j$ and pick the one that minimizes $d(f_i, b_j)$. The time complexity of this algorithm is $\Theta(mn)$.

4.2 Speed-Up with Provable Error Bound

We now present a more efficient video matching scheme detailed in Algorithm 1 (cf. NEAR-LINEAR). The main intuition here is that rather than iterating through every foreground frame in search of its match in the background frame sequence, we only consider foreground frames in chosen intervals of k. The remaining foreground frames can then be matched based on the two considered foreground frames sandwiching it. To obtain an additional improvement, a similar idea is leveraged in the background video, by making comparisons only in jumps of k background frames.

Due to path smoothness, we know that every foreground frame that got *omitted* in consideration is *close enough* to a chosen foreground frame, in particular, at most $\frac{k}{2}$ steps away. Furthermore, due to capture completeness, there exists a background frame that is sufficiently close to each chosen frame. Even though we may have skipped the ideal match, the crucial observation here is that we can apply smoothness on the background as well. This intuition can indeed be formalized concretely. Related ideas that enhance this algorithm are reserved for the next subsection. The following theorem formalizes the error bounds.

Theorem 4.2. *Algorithm* NEAR-LINEAR(k) *runs in* $O(\frac{nm}{k^2})$ *time and returns* π *such that every foreground matches to a background with additive error of at most* $k\delta$*, implying* $C(\pi) \leq C(\pi^*) + k\delta \leq \epsilon + k\delta$.

Note that we may choose k as any value between 1 and $\min\{n, m\}$. This allows us to trade-off between efficiency and accuracy parametrically. For example, if we choose $k \approx \sqrt{m}$, we get a near-linear running time. We state a corollary for an alternative parameter choice that admits the k^2 speed-up while guaranteeing near-optimality.

Corollary 4.3. *For* $k = \frac{\epsilon}{\delta}$, NEAR-LINEAR($k$) *guarantees an* $O(\epsilon)$ *cost, and yet runs* k^2 *times faster.*

The effect of varying k on the algorithm's accuracy is empirically analyzed in a systematic manner in § 6.

4.3 Clustered Matches

We now show that given any foreground frame, if there are *strong matches* (defined below) in the background, they tend to be in close temporal proximity of each other. This understanding allows for a more efficient matching algorithm that can exploit a *local* constraint. In this section we show the theoretical result based on a realistic randomized model that we introduce, and § 6 corroborates our findings through controlled simulation experiments.

Recall that a background frame b corresponds to a strong match for f if the distance $d(f, b)$ is *small*. Let us formalize this notion more crisply here.

Strong Match: We define b to be a strong match for f if $d(f, b) \leq \Psi$. Moreover, given a strong match $\{f, b\}$, we model the measure of distance $d(f, b)$ as a sample from the uniform distribution $[0, \Psi]$.

Before presenting the main result of this subsection, we first state a few prerequisite results.

Lemma 4.4. *Suppose that* b_i *is a strong match for a foreground frame* f*. Then, the probability that* $b_{i \pm \gamma}$ *is also a strong match for* f *is at least* $\frac{\Psi - \gamma\delta}{\Psi}$.

Lemma 4.5. *Given a strong match between* f *and* b_i*, the expected number of strong matches for* f *in the background sequence range* $[b_{i-\gamma}, b_{i+\gamma}]$ *is at least* $2(\gamma + 1)(1 - \frac{\delta\gamma}{2\Psi}) - 1$.

Recall that capture completeness only gives an upper bound guarantee even for the *best* match, specifically a distance of at most ϵ. We therefore consider *strong* matches as having a weaker upper bound guarantee, *i.e.*, $\Psi \approx c\epsilon$ (i.e. $\frac{\Psi}{\delta} \gg c$). In such a situation, we wish to interpret Lemma 4.5 with reasonable parameter values. We present one such choice of parameters in the following theorem and then proceed to describe the algorithm that would guarantee efficiency in terms of a speed-up similar to k^2 as before, without the added cost of missing out on strong matches.

Theorem 4.6. *Given a strong match, with a window size* $\gamma = \frac{\Psi}{\delta}$*, the expected number of strong matches in the entire window of size* $(2\gamma + 1)$ *is at least* γ*.*

This leads to an algorithm that suggests making γ skips in the background frames and whenever encountering a strong match, adopting a local search approach. Theorem 4.6 therefore suggests finding several strong matches without performing exhaustive search. This procedure is listed for one foreground frame in Algorithm 2, and can be extended $\forall f \in \bar{f}$ with strong background matches. For foreground frames with no strong background matches, we propose to rely on foreground tracking (see § 5.6 for details).

5. COMPUTATIONAL FRAMEWORK

We now present a foreground extraction framework that exploits our frame matching algorithm as explained in § 4. The main steps of our framework are illustrated in Figure 2. There are two important insights to be noted here:

• While our scene might have depth, we show that using projective transform followed by non-rigid frame warping generally works sufficiently well for aligning the background and foreground frames.

• Furthermore, the warping error obtained for a foreground frame varies over different matched background frames. This allows us to integrate over multiple noisy warping hypotheses to obtain a robust inference about the foreground.

Recall that approximating scene geometry without using the complete 3-D scene structure has previously been explored [24] [22]. The novelty of our framework however lies in applying this idea to obtain *multiple* noisy estimates of foreground, and to then *fuse* these multiple hypotheses to obtain a final robust inference. Our results show that combining these two steps is crucial for efficient background subtraction of videos tracing substantially different camera paths. We now explain the steps of our framework.

5.1 Key-Point Detection & Frame Matching

For each foreground and background frame, we compute key-points descriptors using the SURF pipeline [3], and find matches between them using the RANSAC algorithm [9].

We define the distance between a pair of background and foreground frames, $d(f, b)$, as the reciprocal of the number of key-points matches found between them[1], and search for corresponding background and foreground frames with strong matches using our proposed method explained in § 4.

5.2 Perspective Transformation

For each foreground frame f, given a strong match $b \in \bar{b}$, we compute a perspective transform $\zeta_{f,b}$ between them [13] such that $b = \zeta_{f,b} f$. Using $\zeta_{f,b}$, we warp b to align it with the given foreground frame [10], and generate the transformed background frame $b^{'}$, such that $b^{'} = \zeta_{f,b}^{-1} b$.

5.3 Non-Rigid Transformation

Using perspective transformation, most parts of f and $b^{'}$ get well-aligned with each other. However, there still may be parts where this alignment is not accurate. This disparity is possible because of the planar assumption that perspective transformation makes regarding the scene. However, for our problem this assumption is not necessarily true. To this end, we first compute the optical flow between f and $b^{'}$ [8]. We then apply the computed flow-fields to warp $b^{'}$ onto f, to find the non-rigidly transformed background frame $b^{''}$ [10].

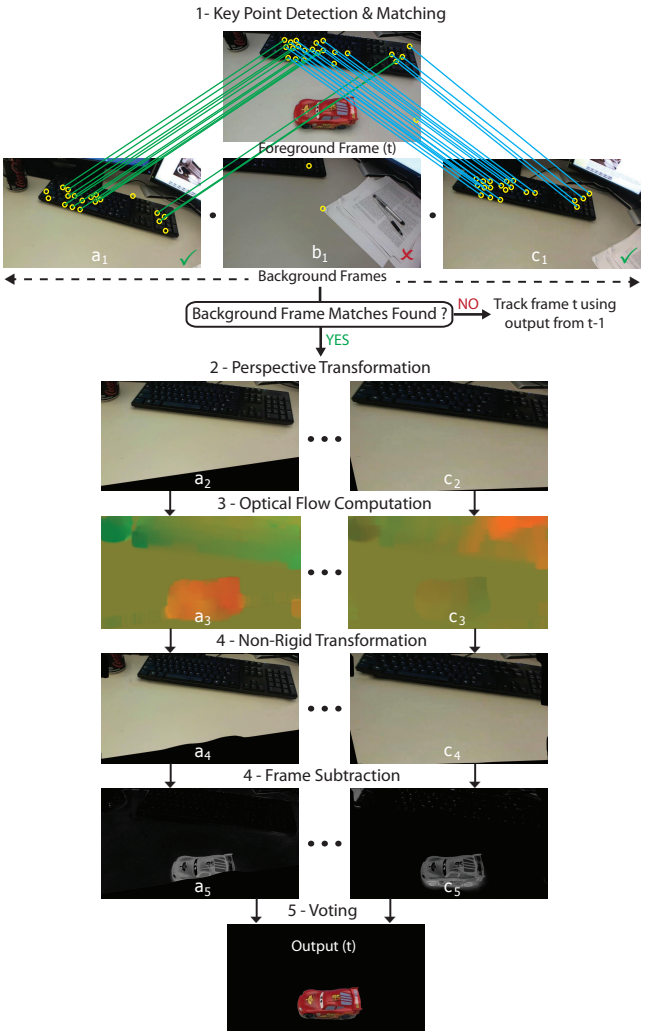

Figure 2: Illustration of the different steps in our proposed computational framework.

5.4 Frame Subtraction

For each pair of f and $b^{''}$, we perform a per-pixel frame subtraction between their corresponding color channels[2]. The subtracted pixels greater than a certain threshold τ_s are classified as belonging to the foreground scene, and used to obtain the foreground mask $\mathcal{F}$, *i.e.*

$$\mathcal{F}(x,y) = \begin{cases} 1 & \text{if } f(x,y) - b^{''}(x,y) > \tau_s \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

[1]While there can be several measures of frame distance, the number of key-point matches has been shown to be a reasonable choice [18].

[2]While color as a measure of pixel difference worked well for us, one could also use more robust pixel features instead *e.g.*, dense SURF [3] *etc.*

Algorithm 3 VOTING-SCHEME($\bar{f}, \bar{b}$)

Input: $f \in \bar{f}$, and a set of strong match background frames $b_s(f) \subseteq \{b_1, b_2, \ldots, b_m\}$. Threshold parameter t.
Output: Background subtracted image corresponding to f.

1: Let $|b_s(f)| = r$.
2: For every pixel that is labeled as foreground on at least τ_v strong matches, mark it as foreground in the result.
3: All other pixels are labeled as background.

5.5 Per-Pixel Voting

The frame subtraction operation gives multiple candidate hypotheses $\mathcal{F}^i(x, y)$ characterizing whether a pixel in the foreground frame belongs to the foreground scene or not. For each of the foreground frame pixels we count the number of times it is classified as belonging to the foreground scene. If this fraction is more than a certain threshold τ_v, it is finally classified as a foreground scene pixel $\mathcal{R}(x, y)$, *i.e.*,

$$\mathcal{R}(x,y) = \begin{cases} 1 & \text{if } \text{avg}_i(\mathcal{F}^i(x,y)) > \tau_v \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

This procedure is listed in Algorithm 3 (cf. VOTING-SCHEME). Note that the likelihood of classifying a pixel correctly increases exponentially with the increase in the number of strong matches. Formally, assume that for any pixel $p \in \mathcal{F}$, we have the following two properties:

- If the pixel is from foreground, then it is reflected as so with probability $\geq p_1$, and
- If the pixel is from background, then it is reflected as so with probability $\geq p_2$ where $p_1 + p_2 > 1$.

Then the following lower-bound holds on the probability of correctly classifying p.

Theorem 5.1. *Given r strong matches, and a voting threshold t in* VOTING-SCHEME *in the open interval $(1 - p_2, p_1)$, the probability that any pixel is classified correctly is at least $1 - e^{-O(r)}$.*

Intuitively, Theorem 5.1 suggests that for each foreground frame, we only need a few background frames with strong matches to perform foreground extraction robustly.

5.6 Selective Foreground Tracking

There might be foreground frames for which there are no strong background matches to be found. For such frames, we propose to track the latest detected foreground from their previous frame that did have strong background matches. For this work, we found Mean-Shift tracking [7] to be an efficeint choice for tracking the foreground when needed.

Since for any foreground frame, tracking the foreground pixels is only a constant time operation, it does not effect the overall complexity of the algorithms described in § 4.

6. SIMULATION ANALYSIS

Since the paths traced by handheld cameras can be unconstrained, it is important to analyze how the accuracy of our framework varies as a function of the difference between the cameras paths traced by the background and foreground videos. We now present a set of controlled simulation experiments to analyze this question in a systematic manner.

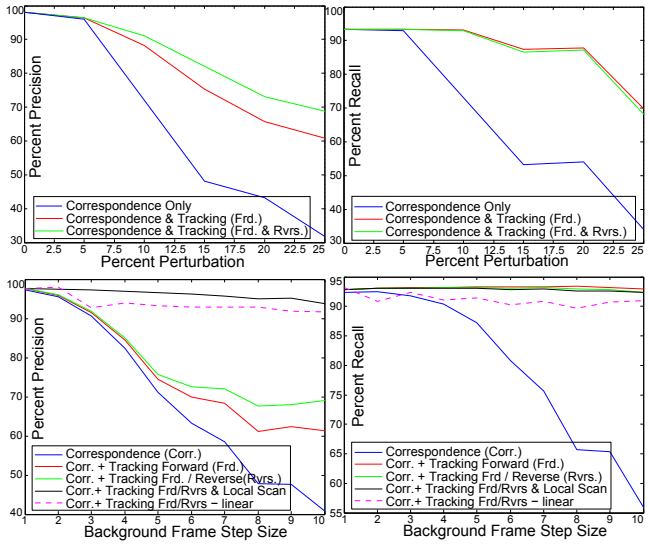

Figure 3: Precision (top left) and Recall (top right) plots with increasing amounts of camera path perturbation. Blue graph shows results obtained by only using frame matching (corr.), while green and red graphs represent also using forward (frd.) and reverse (rvrs.) tracking of foreground pixels in frames that do not have strong matches in the background video. Precision (bottom left) and Recall (bottom right) are shown as a function of background frame step-size. The black and purple graphs use local search in the background, and both background and foreground videos.

6.1 Video Data Generation

We used a 3-D animation software (Blender) to design a desktop scene (Figure 1) where we modeled the background camera path as a circle parameterized by a set of key points. To generate the foreground camera path, we randomly perturbed these key-points by different amounts, and fitted a spline curve to them. We quantified the added perturbation added as a fraction of its radius. For each perturbation amount, we generated 10 sets of foreground paths.

6.2 Simulation Experiments

We ran three main experiments on our generated data.

Experiment 1: In this experiment (Figure 3, top row), we analyzed our precision and recall as a function of foreground path perturbation[3]. Here, we searched for strong matches in a naive manner (§ 4.1).

Experiment 2: For our second experiment, we analyzed how strong background frame matches are temporally distributed. We performed this experiment for different amounts of foreground path perturbations (Figure 4).

Experiment 3: Finally, for our third experiment (Figure 3, bottom row), we analyzed for a given path perturbation, how does our performance vary as we search for strong background frame matches in a more greedy way (§ 4.2).

[3]As we know the ground truth, we can compute the precision and recall based on their standard definitions.

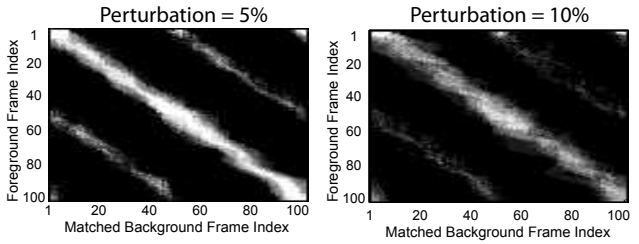

Figure 4: For each foreground frame, we show how likely a background frame is to be a *strong match*. Here white implies high likelihood. Results for 5% & 10% perturbation are shown. Greater perturbation amounts follow similar trends.

6.3 Main Findings

We now report our main findings from these experiments.

6.3.1 Matching Alone is Not Enough

When purely relying on finding a strong matching between the foreground and the background frames, the performance of our framework degrades steeply (blue plots in Figure 3, top row). This is because for larger disparity between the background and foreground camera paths, there may be no strong matches available, *i.e.*, capture completeness (§ 3.1.2) may not be true for all foreground frames. However, as can be seen from the red plots in the top row of Figure 3, adaptively switching between background subtraction of registered frames, and tracking the foreground pixels based on their latest available location information substantially improves the performance.

6.3.2 Temporal Smoothing Improves Accuracy

Performing foreground tracking for frames with no matching background can be improved by combining a forward and a reverse tracking pass (Figure 3, top row). This is because for the foreground frames for which we do have strong background matches, the noisy pixels appear and disappear erratically, while the foreground pixels are much more consistent over frames. Therefore, fusing information of foreground pixels obtained by tracking them in the forward and reverse direction results in accuracy improvement at only a constant additional cost in time complexity.

6.3.3 Matches are Temporally Clustered

We analyzed how strong background frame matches are temporally distributed in the background video. This analysis corroborates our theoretical result of Theorem 4.6 which suggested that several strong matches are likely to be clustered together. This can be observed from Figure 4, where light regions show pairs of foreground and background frames that form a strong match (averaged over 10 independent trials). As can be seen, light regions cluster together, lending empirical evidence to our theoretical result.

6.3.4 Selective Coarse-to-Fine Frame Search Works

To find strong matches for a foreground frame, it is an efficient and accurate strategy to first sparsely search a well-distributed set of background frames, followed by a more dense search around background frames that match well with the given foreground frame. This is in accordance with Algorithm 2, and is empirically corroborated by the results shown in Figure 3 (bottom row).

As can be seen, with the forward and reverse tracking, and the local search principle incorporated (the dashed purple graph), the algorithm's performance (both for precision and recall) is nearly unaffected even as the step size is increased. Specifically, the precision and recall are both above 90% for step sizes all the way up to $k = 10$. Note that as the step size increases, it significantly outperforms all approaches that do not leverage local search. Further, it comes quite close to the black graph (the overall best) despite being an additional k times faster than it. Note that for these experiments, the number of video frames was 100 each. This implies that our framework can achieve greater than 90% accuracy while taking time only linear in the lengths of the background and foreground videos.

7. EXPERIMENTS & RESULTS

We now present results of our framework tested on real videos taken by multiple users.

7.1 Explored Test Scenarios

We performed a set of experiments designed to test different challenges for handheld cameras. Sample video frames and output from for these are presented in Figure 5. The full videos are provided in the supplementary material.

Foreground Object Geometry: The first five results in figure 5 are presented in the order of increasing foreground complexity. As shown, our framework performs well for a range of object geometries including mostly convex (result 1) to quite non-convex (result 5).

Indoors Versus Outdoors: In Figure 5, results 1 to 4 were captured indoors, while result 5 was captured outdoors in daylight. Even within the indoor environment the auto white-balancing feature of mobile cameras can result in substantial variation in the global illumination of the video. Our framework is able to handle these variations quite well.

Multiple Foreground Objects: Our framework is agnostic to the number of foreground objects present in the video. This is shown in result 6 of Figure 5, where we were able to successfully extract all three foreground objects.

Moving Foreground Objects: We captured videos for both simple and articulated moving objects, shown in the last two results of Figure 5. Not only do we preform well for the limited movement of the table fan, we were also able to extract the articulated human motion. This result is a first step for automatic green-screening of user-generated videos.

7.2 Comparative Results

We compared our algorithm with five alternate foreground extraction methods. These methods are explained below.

- In the first comparison (blue plot), we used a Gaussian Mixture model (GMM) to learn the global background appearance which was used for background subtraction [25].
- In the second comparison (green plot), we manually provided a foreground mask for the first frame of the foreground video to learn a foreground GMM. This model was used to subtract background from all the foreground frames.
- Our third comparison (green plot) was with the previously proposed approach of video matching [21], and highlights its limitations due to its strong assumptions about the foreground and background camera paths being the same.
- Our fourth comparison (orange plot) was with a perceptual

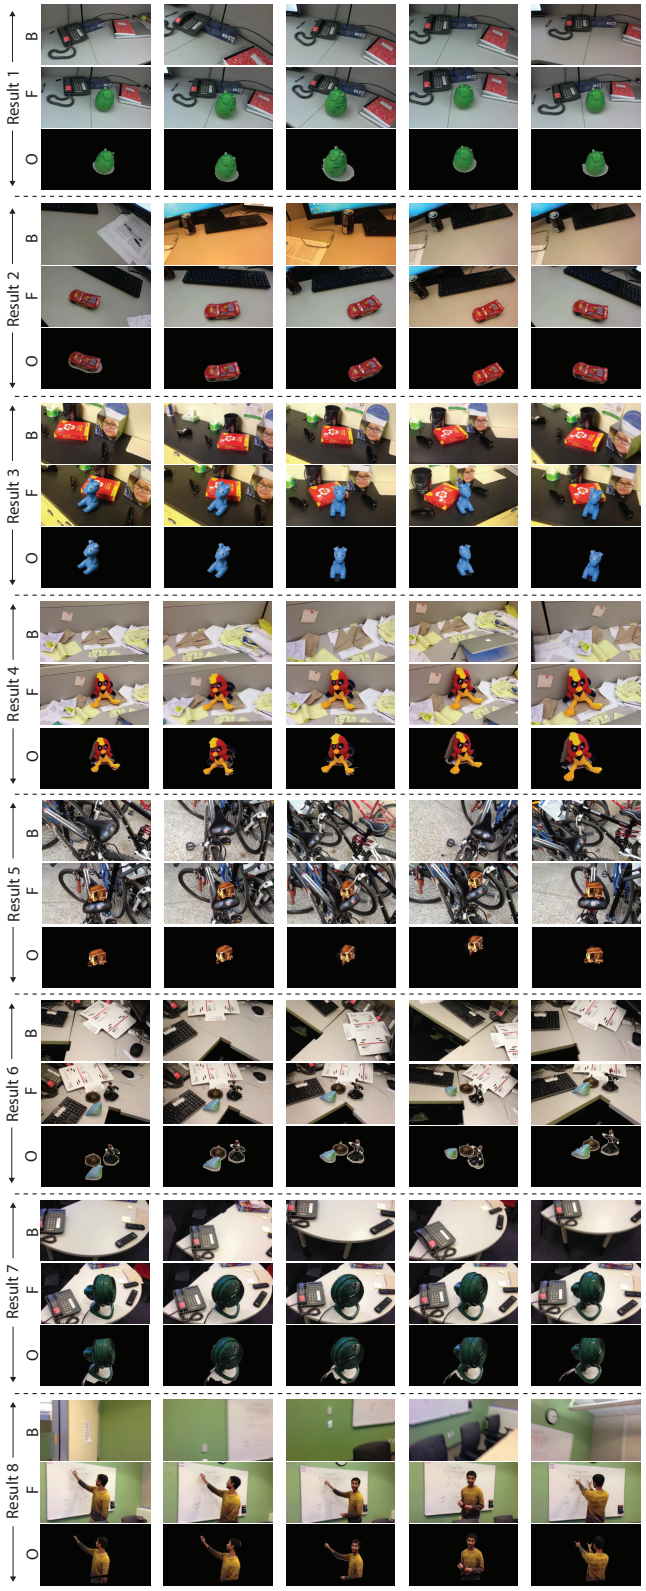

Figure 5: Results capturing various experiment parameters are presented. Here B, F, and O denote background, foreground, and output videos.

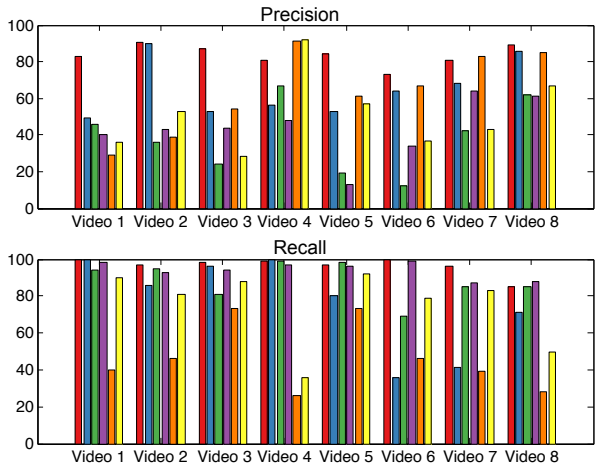

Figure 6: Precision and Recall plots of our proposed method (red) and five comparative approaches.

saliency based approach for the foreground frames [16].
• Our fifth and final comparison (yellow plot) was with a multi-frame co-segmentation based approach [19].

7.2.1 *Precision & Recall*

We manually extracted the foreground object(s) in every 10^{th} frame of the test videos presented in Figure 5. Using this, we computed the number of true/false positives and negatives. We report the average precision and recall for our method in Figure 6 (red plot). As shown, our approach gives better recall as well as comprehensively outperforms previous approaches on precision.

7.2.2 *Time Comparisons*

Clock times taken by three of the five considered algorithms on the test videos are shown in Table 1. Note that the remaining two algorithms from Figure 6 are derivatives of the three considered methods in Table 1, with times similar to the first two rows of the table. As can be observed, our proposed method takes the least amount of processing time while generating superior results compared to all the considered alternatives.

8. CONCLUSIONS & FUTURE WORK

In this paper, we identified a novel instance of the background subtraction problem focusing on extracting near-field foreground objects captured using handheld devices. We posed our problem as efficient matching of foreground and background videos tracing significantly different camera trajectories. We proposed an approximate algorithm for this problem with provable theoretical guarantees that

	V_1	V_2	V_3	V_4	V_5	V_6	V_7	V_8
A_1	41.7	11.8	19.1	5.8	83.6	42.5	18.6	47.1
A_2	10.6	6.8	11.7	4.7	5.9	5.6	8.7	10.6
A_3	**8.1**	**4.3**	**8.7**	**2.9**	**5.2**	**5.3**	**5.7**	**9.5**

Table 1: Algorithms' clock times in minutes. Here A1, A2, and A3 represent GMMs based modeling, saliency based extraction, and proposed method.

admit a significant computational speed-up. Using insights from our theoretical model, we proposed an end-to-end computational framework, and presented controlled simulation experiments to thoroughly analyze its object-extraction accuracy. Moreover, we presented results of our framework on multiple hand-held videos, and compared them with five alternate approaches. Going forward, we plan to also learn the shape and appearance of the extracted foreground object(s). We expect this to allow us to improve our accuracy further, while requiring even fewer matches.

APPENDIX

A. THEORETICAL ANALYSES

In the following we present the detailed proofs for the theorems and lemmas we presented in our paper.

Lemma 3.1 $C(\pi^*) \leq \epsilon$.

Proof: Follows from the path smoothness assumption.

Theorem 3.2: Algorithm NEAR-LINEAR(k) runs in $O(\frac{nm}{k^2})$ time and returns π such that every foreground matches to a background with additive error of at most $k\delta$, i.e. $C(\pi) \leq C(\pi^*) + k\delta \leq \epsilon + k\delta$.

Proof: The proof follows in the following two steps.

Step 1: We first show the bound $\forall\ f_i$ where $i \in S_i^k$. Let $j_{opt} = \arg\min d(f_i, b_{j_{opt}})$. In algorithm NEAR-LINEAR(k), let the returned value for $\pi(i) = j_{alg}$. Let $near(j_{opt}) = \arg\min_{t \in S_j^k} |t - j_{opt}|$. It follows that $|near(j_{opt}) - j_{opt}|$ is at most $k/2$, since S_j^k has all multiples of k as well as 1 and m. Therefore by triangular inequality, $d(f_i, b_{near(j_{opt})}) \leq d(f_i, b_{j_{opt}}) + d(b_{j_{opt}}, b_{near(j_{opt})})$. This is at most $d(f_i, b_{j_{opt}} + \delta \cdot |near(j_{opt}) - j_{opt}|$ by the continuous capture assumption and the triangle inequality (since $d(b_x, b_{x+1}) \leq \delta$ for all x). It follows that $d(f_i, b_{near(j_{opt})}) \leq d(f_i, b_{j_{opt}}) + k\delta/2$. Notice that $d(f_i, b_{j_{alg}}) \leq d(f_i, b_{near(j_{opt})})$, because $near(j_{opt}) \in S_j^k$ and the algorithm picked $\pi(i) = j_{alg}$ as $\arg\min_{j \in S_j^k} d(f_i, b_j)$. Combining, we get $d(f_i, b_{j_{alg}}) \leq d(f_i, b_{j_{opt}}) + k\delta/2$. Furthermore, we know by the complete capture assumption, that $d(f_i, b_{j_{opt}}) \leq \epsilon$. It follows that $d(f_i, b_{j_{alg}}) \leq \epsilon + k\delta/2$.

Step 2: We now consider the case when $i \notin S_i^k$. Let $near(i) = \arg\min_{t \in S_i^k} |t - i|$. It follows that $|i - near(i)| \leq k/2$, because S_i^k contains all multiples of k between 1 and n, and also includes $\{1\}$ $\{n\}$. As in step 1, let the returned value for $\pi(i)$ be j_{alg}. Therefore, by triangle inequality $d(f_i, b_{j_{alg}}) \leq d(f_{near(i)}, b_{j_{alg}}) + \delta \cdot |i - near(i)|$. This is because using the triangle inequality and the continuous capture assumption, $d(b_x, b_y) \leq \delta \cdot |x - y|$. Since j_{alg} is picked as $\pi(near(i)$, it follows that $d(f_{near(i)}), b_{j_{alg}}) \leq d(f_i, b_{j_{opt}}) + k/\delta 2$. Combining, we get $d(f_i, b_{j_{alg}}) \leq d(f_i, b_{j_{opt}}) + k\delta/2 + \delta \cdot |i - near(i)| \leq d(f_i, b_{j_{opt}}) + k\delta/2 + k\delta/2 = d(f_i, b_{j_{opt}}) + k\delta$. As before, $d(f_i, b_{j_{opt}}) \leq \epsilon$ by the complete capture assumption. It follows that $d(f_i, b_{j_{alg}}) \leq \epsilon + k\delta$. Averaging over all i, this implies that $C(\pi) \leq C(\pi^*) + k\delta \leq \epsilon + k\delta$.

Making use of the above two cases, we now show the time complexity analysis of the algorithm NEAR-LINEAR(k). Notice that $|S_i^k| \leq \lfloor n/k \rfloor + 2$ since S_i^k only has multiples of k in $[1, n]$ and $1, n$. Similarly, $|S_j^k| \leq \lfloor m/k \rfloor + 2$. Since the algorithm only computes distances between pairs f_i and b_j where $i \in S_i^k$ and $j \in S_j^k$, and the other computations are just finding $near(i)$ for $i \notin S_i^k$, the total computation time is $O(|S_i^k| \cdot |S_j^j|)$, which is $O(mn/k^2)$. This completes the proof.

Corollary 3.3 For $k = \frac{\epsilon}{\delta}$, NEAR-LINEAR($k$) guarantees an $O(\epsilon)$ cost as the optimal algorithm, but runs k^2 times faster.

Proof: Recall that the optimal guarantees only a cost of ϵ. Now, using Theorem 3.2 plugging in $k = \frac{\epsilon}{\delta}$ gives a cost for algorithm NEAR-LINEAR equal to $\epsilon + k\delta = 2\epsilon$. Further, the algorithm running time is mn/k^2 or $\frac{mn\delta^2}{\epsilon^2}$. Notice that for $\epsilon \gg \delta$ this is a significant improvement on running time at almost no cost on the solution guarantee.

Lemma 3.4 Suppose that b_i is a good match for a foreground frame f. Then, the probability that $b_{i \pm \gamma}$ is also a good match for f is at least $\frac{\Psi - \gamma\delta}{\Psi}$.

Proof: Since b_i is good match for f, we know that $d(f, b_i)$ is at most Ψ and further is a random variable uniformly distributed in $[0, \Psi]$. Recall by the continuous capture characteristic and the triangle inequality of the distance metric, we know that $d(b_i, b_{i+\gamma}) \leq \gamma\delta$. So applying triangle inequality again, we have $d(f, b_{i+\gamma}) \leq d(f, b_i) + d(b_i, b_{i+\gamma})$.

Therefore, we know that $d(f, b_{i+\gamma})$ is a random variable that is dominated from above by a random variable that is uniformly distributed in $[\gamma\delta, \Psi + \gamma\delta]$. The probability that such a sample is at most Ψ is at least $\frac{\max\{\Psi - \gamma\delta, 0\}}{\Psi}$ (we drop the max in the lemma statement as if the quantity becomes negative, the lemma is vacuously true). Therefore, this is the probability that $b_{i+\gamma}$ is also a good match for f. The exact same proof carries over for $b_{i-\gamma}$ as well.

Lemma 3.5 Given a good match between f and b_i, the expected number of good matches for f in the background sequence range $[b_{i-\gamma}, b_{i+\gamma}]$ is at least $2(\gamma + 1)(1 - \frac{\delta\gamma}{2\Psi}) - 1$.

Proof: From Lemma 3.4, we know that if b_i is a good match for f, then the probability that b_j is a good match for f is at least $\frac{\Psi - \delta|j - i|}{\Psi}$. Now by using linearity of expectation for all j in range $[i - \gamma, i + \gamma]$, we get that the number of expected matches for f in this range for b is at least

$$\begin{aligned}
&\textstyle\sum_{j=i-\gamma}^{i+\gamma} \frac{\Psi - \delta|j-i|}{\Psi} \\
&= 1 + 2 \cdot \textstyle\sum_{j=i+1}^{i+\gamma} \frac{\Psi - \delta(j-i)}{\Psi} \\
&= 1 + 2(\gamma - \tfrac{\delta}{\Psi}(1 + 2 + \ldots + \gamma)) \\
&= 1 + 2(\gamma - \tfrac{\delta}{\Psi} \cdot \tfrac{\gamma(\gamma+1)}{2}) \\
&= (2\gamma + 1) - \tfrac{\delta\gamma(\gamma+1)}{\Psi} \\
&= 2(\gamma + 1)(1 - \tfrac{\delta\gamma}{2\Psi}) - 1.
\end{aligned}$$

Theorem 3.6 Given a good match, with a window size $\gamma = \frac{\Psi}{\delta}$, the expected number of good matches in the entire window of size $(2\gamma + 1)$ is at least γ.

Proof: From Lemma 3.5, given a good match b_i, the expected number of good matches between $b_{i-\gamma}$ and $b_{i+\gamma}$ is at least $2(\gamma+1)(1 - \frac{\delta\gamma}{2\Psi}) - 1$. Using $\gamma = \frac{\Psi}{\delta}$ the theorem follows.

Theorem 4.1 Given r good matches, and a chosen voting threshold t in VOTING-SCHEME, the probability that any pixel is classified correctly is at least $1 - e^{-O(r)}$.

Proof: Given independent identically distributed random variables X_i that take on values 0 or 1, define $X = \sum_i X_i$ and let μ be the expectation of X. Then, Chernoff bound states that $P(|X - \mu| \geq \alpha \cdot \mu) \leq e^{(-\frac{\mu\alpha^2}{2})}$. If a pixel is originally an object pixel, then in the observed matches, it has probability $\geq p_1$. Therefore, if each X_i is defined as the random variable for whether in a good match i, the pixel gets marked as an object, then we have $\mu = E(X) \geq p_1 r$. So the probability that it gets marked as an object in fewer than tr pixels is $P(|X - p_1 r| \geq \frac{|t-p_1|}{p_1}(p_1 r)) \leq e^{-p_1 r(\frac{t-p_1}{p_1})^2/2}$. This is $\leq e^{-(\frac{(t-p_1)^2}{2p_1})\cdot r}$. Now suppose that a pixel was originally a non-object pixel, then in the observed matches, it is classified as a non-pixel with probability $\geq p_2$. Therefore if each X_i is defined as the random variable for whether in a good match i, the pixel is incorrectly marked as an object, then we have $\mu = E(X) \leq (1 - p_2)r = (p_1 - \beta)r$ for some constant $\beta = (p_1 + p_2 - 1) > 0$. Using similar argument as the previous paragraph, we get the probability that the pixel gets marked incorrectly as an object is at most $e^{-(\frac{(t-(p_1-\beta))^2}{2(p_1-\beta)})\cdot r}$. Setting t to any value between $(p_1 - \beta)$ and p_1, for example $t = p_1 - \frac{\beta}{2}$, we have both quantities $(t - p_1)^2$ and $(t - (p_1 - \beta))^2$ are constants *strictly* bigger than 0. It follows that the probability of error in both cases is at most $e^{-(\frac{\beta^2}{8p_1})\cdot r}$ which is $e^{-O(r)}$. The theorem follows.

B. RESULT VISUALIZATIONS

For the sake of visual comparison, we show the results of considered approaches on a subset of frames for six of the considered test videos in Figure 7. The results for the remaining test videos show similar visual trends.

C. REFERENCES

[1] O. Aghazadeh, J. Sullivan, and S. Carlsson. Novelty detection from an ego-centric perspective. *IEEE CVPR*, pages 3297–3304, 2011.

[2] O. Aghazadeh, J. Sullivan, and S. Carlsson. Multi view registration for novelty/background separation. *CVPR*, 2012.

[3] H. Bay, T. Tuytelaars, and L. Van Gool. Surf: Speeded up robust features. *Computer Vision–ECCV 2006*, pages 404–417, 2006.

[4] A. Bevilacqua and P. Azzari. High-quality real time motion detection using ptz cameras. *AVSS'06*, 2006.

[5] R. Carceroni, F. Pádua, G. Santos, and K. Kutulakos. Linear sequence-to-sequence alignment. *CVPR*, 2004.

[6] Y. Caspi and M. Irani. Spatio-temporal alignment of sequences. *PAMI*, 2002.

[7] D. Comaniciu, V. Ramesh, and P. Meer. Real-time tracking of non-rigid objects using mean shift. *CVPR*, 2:142–149, 2000.

[8] G. Farneback. Fast and accurate motion estimation using orientation tensors and parametric motion models. In *ICPR*, 2000.

[9] M. Fischler and R. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of ACM*, 24(6):381–395, 1981.

[10] C. Glasbey and K. Mardia. A review of image-warping methods. *JAS*, 25(2):155–171, 1998.

[11] M. Granados, K. I. Kim, J. Tompkin, J. Kautz, and C. Theobalt. Background inpainting for videos with dynamic objects and a free-moving camera. In *ECCV (1)*, pages 682–695, 2012.

[12] J. Harel, C. Koch, and P. Perona. Graph-based visual saliency. In *NIPS*, pages 545–552, 2006.

[13] R. Hartley and A. Zisserman. Multiple view geometry in computer vision. 2, 2000.

[14] E. Hayman and J. Eklundh. Statistical background subtraction for a mobile observer. *CVPR*, 2003.

[15] M. Isard and J. MacCormick. Bramble: A bayesian multiple-blob tracker. *ICCV*, 2:34–41, 2001.

[16] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *PAMI*, 20(11), 1998.

[17] A. Kowdle, S. N. Sinha, and R. Szeliski. Multiple view object cosegmentation using appearance and stereo cues. In *ECCV*, 2012.

[18] C. Liu, J. Yuen, and A. Torralba. Sift flow: Dense correspondence across scenes and its applications. *PAMI*, 2011.

[19] F. Meng, H. Li, G. Liu, and K. Ngan. Object co-segmentation based on shortest path algorithm and saliency model. 2010.

[20] K. Papoutsakis and A. Argyros. Object tracking and segmentation in a closed loop. *AVC*, 2010.

[21] P. Sand, S. Teller, et al. Video matching. *ACM Transactions on Graphics*, 23(3):592–599, 2004.

[22] S. M. Seitz and C. R. Dyer. Complete scene structure from four point correspondences. In *ICCV*, 1995.

[23] Y. Sheikh, O. Javed, and T. Kanade. Background subtraction for freely moving cameras. *CVPR*, 2009.

[24] H. Shum and R. Szeliski. Construction and refinement of panoramic mosaics with global and local alignment. In *ICCV*, 1998.

[25] C. Stauffer and W. Grimson. Adaptive background mixture models for real-time tracking. *CVPR*, 2, 1999.

[26] J. Sun, W. Zhang, X. Tang, and H. Shum. Background cut. *IEEE ECCV*, pages 628–641, 2006.

[27] P. Torr, R. Szeliski, and P. Anandan. An integrated bayesian approach to layer extraction from image sequences. *PAMI*, 23, 2001.

[28] K. Toyama, J. Krumm, B. Brumitt, and B. Meyers. Wallflower: Principles and practice of background maintenance. *CVPR*, 1999.

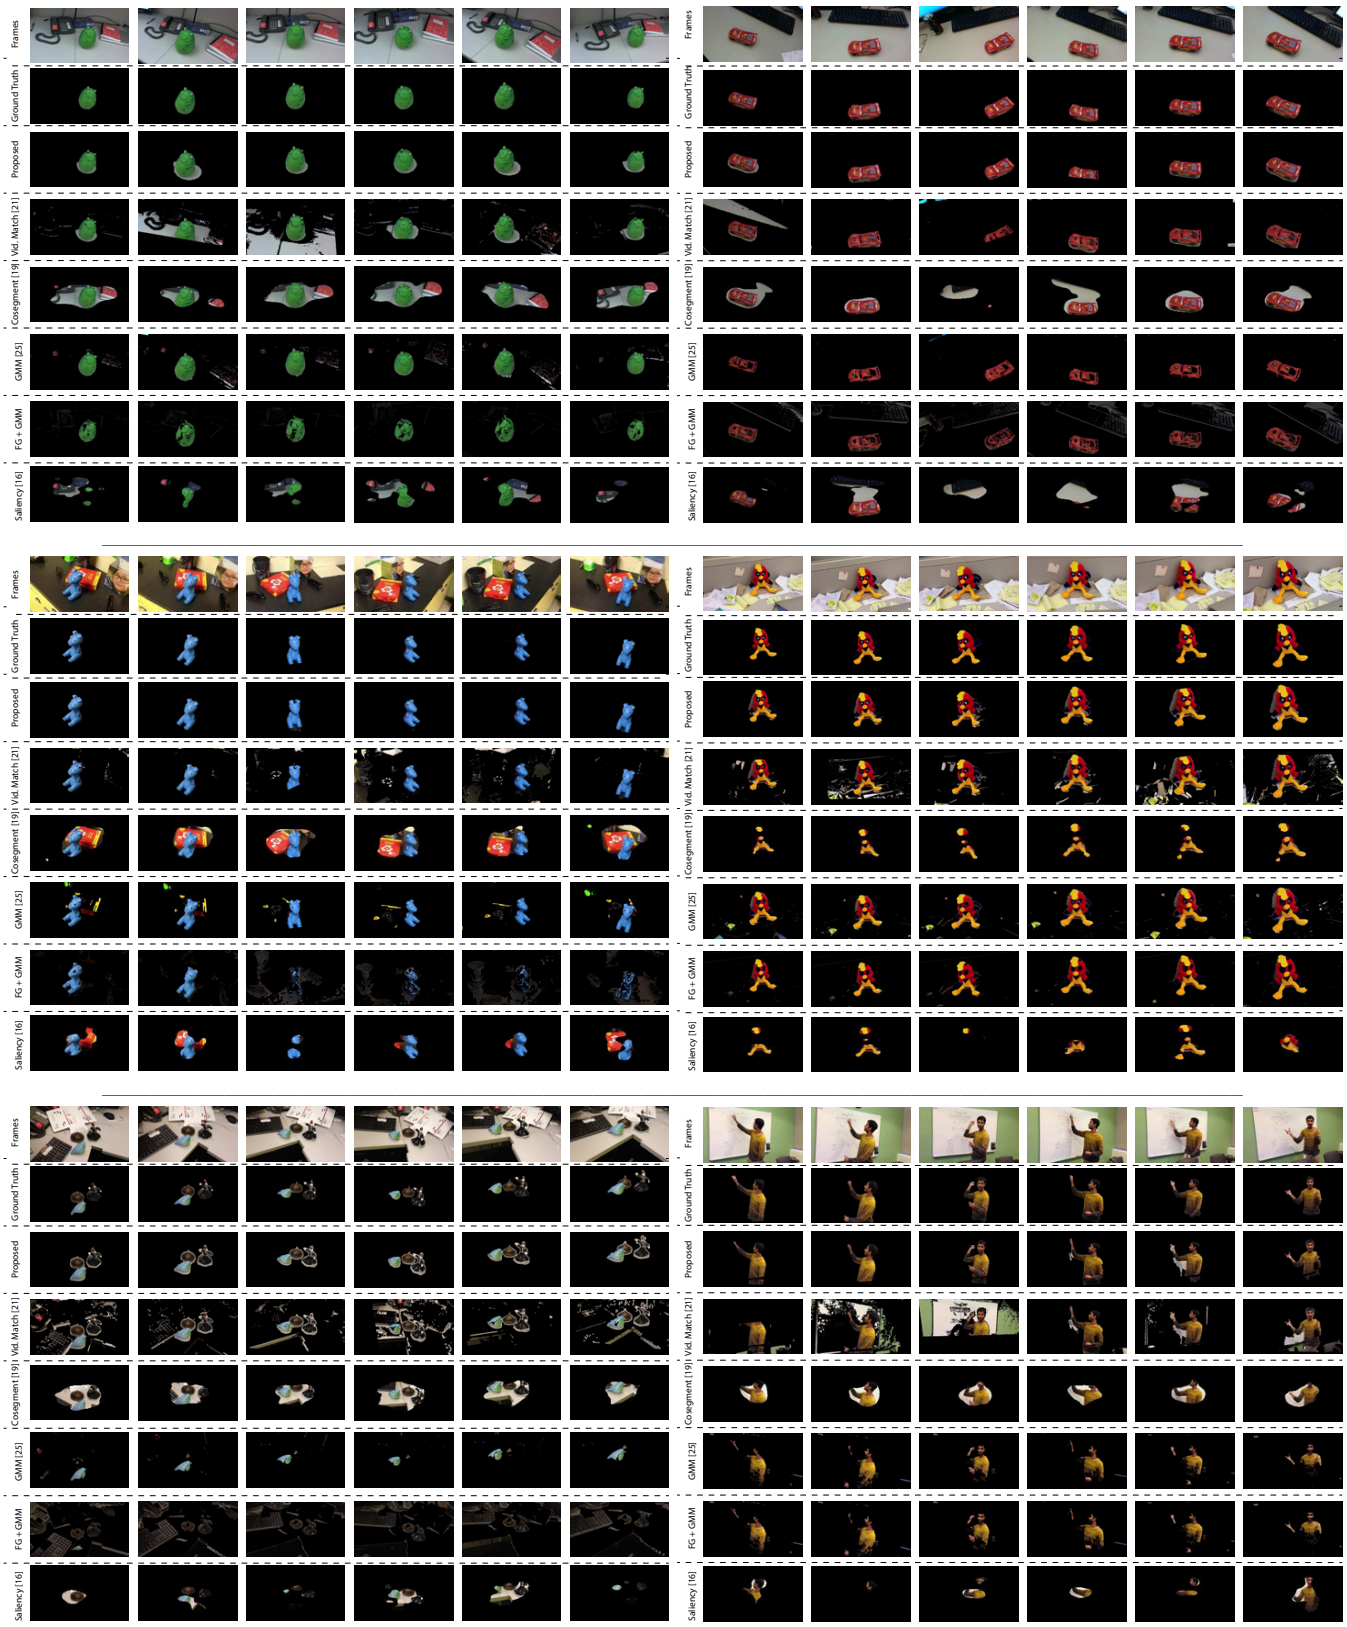

Figure 7: Comparative visualizations of subset of frames for some of the test videos are shown. As can be seen, the quality of the output for the proposed method is very similar to that of the manually extracted ground-truth. Also note that the quality of the proposed method is far superior than any of the other comparative algorithm.